HP recommended configurations for online transaction processing: ProLiant DL980 G7, VMA-series Memory Array (VMA) and Oracle 11gR2 database

When your OLTP application needs extreme performance

Technical white paper

Table of contents

Executive summary

Always a critical part of the IT infrastructure, databases are at the heart of a number of trends in IT:

- Demand for large symmetric multiprocessing (SMP) type scalability in a Linux or Microsoft® Windows® environment
- Complex workloads that demand better response times and increased I/O operations per second (IOPS)
- Oracle applications and data are growing in size and complexity. Some applications running on 32-bit platforms may be reaching the limits of the architecture, specifically regarding the amount of addressable memory.
- Organizations are now able to migrate complex solutions from high-cost proprietary systems to lower-cost Linux or Windows-based systems.
- Organizations are consolidating applications to fewer servers to simplify critical data center operations. Consolidation can reduce cost, management complexity, and physical space requirements in the data center.
- Historically, different computer system components have advanced at different relative rates. Although disk capacity has improved somewhat, disk performance ranks at the bottom with no significant improvement compared to million-fold boosts by other system components such as servers.

Online transaction processing (OLTP) applications such as enterprise resource planning (ERP), supply chain management (SCM), web-based e-commerce systems and even extract, transform and load (ETL) applications can benefit from the performance of a scale-up architecture. More importantly, the HP ProLiant DL980 G7 with HP VMA-series Memory Array solution offers the capacity to consolidate and manage these applications in a single server. Any database application that requires increase in performance including lower user response times, and higher throughput and IOPS is a candidate for these configurations.

Systems such as large e-commerce websites that must respond to spikes in demand from large numbers of users and a high volume of transactions should also see improvement. With these applications, large amounts of memory are required to maintain connection context for every database object opened by a user in addition to storage acceleration for significant I/O performance.

This white paper describes three recommended configurations to support small, medium and large OLTP solutions, capitalizing on the power of the HP ProLiant DL980 G7 server, HP VMA3210 Memory Array appliance, Linux, and Oracle 11gR2 database. The objective of these configurations is to meet or exceed the demanding OLTP performance requirements, reduce the time to purchase and implement the solution as well as to minimize any implementation risks.

The recommended configurations described in this paper include the storage configuration, Oracle database setup, best practices, and a bill of materials.

The configuration information is based on data collected in HP's solution performance lab and proof of concept testing.

Target audience: This white paper is intended for system architects and decision makers who have responsibility for evaluating, purchasing, recommending and/or implementing database applications based on Oracle databases. Previous knowledge of Linux, Oracle database software, HP ProLiant servers, and HP VMA-series Memory Array products would be advantageous but is not required.

This white paper describes testing performed in October 2010 and May 2011.

Introduction

The recommended configurations are based on an online transaction processing (OLTP) workload. The set of transactions used in the workload is representative of a web-based application.

These transactions allow users to check product availability, place an order for products, pay for their order, view the status of their order, and schedule delivery of their order. The most common transactions are placing and paying for an order; these two transactions occur with similar frequency and occur approximately ten times as frequently as other transactions. Checking for product availability and the status of an order are simple, read-only transactions. Placing an order for products, paying for an order, and scheduling delivery of an order include read and write (Insert, Select, and Update) operations. Placing an order is the most complex transaction, involving multiple tables. The recommendations provided are most accurate for applications with transactions similar to those described above. Also these solutions will perform very well for packaged solutions such as Oracle E-Business Suite, PeopleSoft, Siebel, and SAP. We strongly recommend that you work with your local HP Reseller or HP Sales Representative to determine the best solution for you.

The objective of the proposed configurations is to:

1. Provide best-in-class performance – You can experience a dramatic increase in performance, efficiency and reliability with HP servers powered by the massively scalable and highly reliable Intel® Xeon® processor E7 family. Built to handle demanding applications, the Intel Xeon processor E7 family delivers a quantum leap in enterprise computing performance and delivers performance scalability that takes server consolidation to the next level. You'll also get the reliability that you need to run your applications with complete confidence, while maintaining data integrity and minimizing downtime.
2. Reduce the risk to implement database solutions – These recommended configurations provide a balance of CPU, memory and I/O capacity specifically optimized for database workloads. As a result, the likelihood of functional or performance problems in your database deployments is drastically reduced.
3. Reduce the time to implement database solutions on HP servers – By using the configuration details described in this document the time to deploy the database and its application can be reduced.
4. Simplify the purchasing process for well-defined HP server configurations – The recommended configurations may be implemented as is or used as a starting point for a semi-custom configuration. In either case, the expectation is that IT can make quicker and more informed decisions to purchase the optimal configuration for the application.

Solution criteria

Table 1 below shows details of the three workloads in terms of the maximum user count, database size, business transactions/hour, and storage type. These values are the estimated maximum values for the recommended configurations listed in this document. The maximum values are not listed for all possible server configurations. A business transaction has to do with a business function, for example, check status of an order or view an invoice. It is not possible to make a direct comparison between a sustained business transaction volume and transactions per minute (TPM). Business transactions are made up of several individual SQL transactions required to satisfy the business request rather than a total of the individual SQL transactions. Some batch and reporting requirements are essential elements of most OLTP environments. If your environment has extensive batch processing and/or reporting requirements, then HP recommends starting with a proof of concept to test your specific workload before the final configuration is defined. Many other configurations can be created that use additional hardware and support a greater number of transactions.

If you need help with a specific Oracle solution or prefer a solution design or sizing based on your requirements please contact your local HP reseller, HP sales representative, or the HP Oracle Solution Center in your region.

Table 1. Recommended Configuration metrics and storage type options

Number	HP ProLiant Server Model	User Count	Database Size TB	Business OLTP Transactions per Hour	Storage Type
1	DL980 G7 (4p/32c)	6,000	8	35,000,000	Flash
2	DL980 G7 (4p/40c)	13,000	15	70,000,000	Flash
3	DL980 G7 (8p/80c)	26,000	25	150,000,000	Flash

The first configuration using four Intel Xeon E7-2830 processors utilizes direct attached storage (DAS) with two HP VMA controllers and Memory Arrays. The second configuration using four Intel Xeon E7-4870 processors utilizes DAS with four HP VMA Memory Arrays and the third configuration uses eight Intel Xeon E7-4870 processors with eight HP VMA Memory Arrays. Additional information on these VMA products can be found in the *Storage configuration details* section.

The tables below list the supported Intel Xeon processors, memory DIMMs, and PCI expansion slots for the ProLiant DL980 G7 server. This information is included to help determine how the recommended configurations can be modified to support different workload sizes or user combinations.

Table 2. Supported E7 Family Processors

Processor Type Intel Xeon	Cores	Max Cores in a DL980 G7
E7-4870 (30MB Cache, 2.4GHz, 130W, 6.4 GT/s QPI)	10	80
E7-2860 (24MB Cache, 2.26GHz, 130W, 6.4 GT/s QPI)	10	80
E7-2850 (24MB Cache, 2.00GHz, 130W, 6.4 GT/s QPI)	10	80
E7-2830 (24M Cache, 2.13 GHz, 105W, 6.4 GT/s QPI)	8	64
E7-4807 (18M Cache, 1.86GHz, 95W, 4.8 GT/s QPI)	6	48

Note:
The Intel Xeon processor E7 series supports Hyper-Threading (HT). HT is recommended and was tested in our configurations. However it is good practice to test HT with your particular application.

Table 3. Supported Memory DIMMs

Memory Kits	Rank
HP 4GB 1Rx4 PC3-10600R-9	Single
HP 4GB PC3L-10600R-9 512Mx	Single
HP 8GB 2Rx4 PC3-10600R-9	Dual
HP 8GB 2Rx4 PC3L-10600R-9, 512Mx RoHS	Dual
HP 16GB 2Rx4 PC3L-10600R-9	Dual
HP 16GB 4Rx4 PC3-8500R-7	Quad

PC3L = low voltage memory

Table 4 represents the minimum, middle, and maximum memory combinations possible for the 4, 8, and 16 GB memory kits available for the DL980 G7 servers. However, for best performance use dual or quad rank memory DIMMs.

Table 4. Minimum, middle, and maximum memory for 4 and 8 processor configurations

Number of CPUs	Memory Density (GB)	Total Memory Cartridges	Min Memory (GB)	Mid Memory (GB)	Max Memory (GB)
4	4	8	64	128	256
4	8	8	128	256	512
4	16	8	256	512	1024
8	4	16	128	256	512
8	8	16	256	512	1024
8	16	16	512	1024	2048

Note:
Max memory depends on the number of processors configured. Four/eight processor configurations support up to 1TB/2TB of memory, respectively. However Red Hat 5.6 only supports up to 1TB. Red Hat 6.x will expand memory support for Oracle once it is available.

Table 5. VMA supported Expansion Slot Configurations

Expansion Slots
• Standard Main I/O with 5 Gen2 slots: (3) x4 PCI-Express; (2) x8 PCI-Express
• PCIe Option with 6 slots: (1) x4 Gen1 PCI-Express; (1) x4 Gen2 PCI-Express (4) x8 Gen2 PCI-Express
• Low Profile Expansion Option with 5 Gen2 slots; (1) x4 PCI-Express; (4) x8 PCI-Express

Note:
The PCI-X and PCIe Combination Option slots are not to be used for HP VMA-series Memory Arrays.

I/O slots capable of bi-directional x8 PCIe provide the best performance. The DL980 G7 server comes with the Standard Main I/O board with PCIe slots 7-11. Slots 9 and 11 are x8 Gen2 PCIe slots. The PCIe expander option provides additional I/O slots 1-6. Slots 2, 3, 5 and 6 are x8 Gen2 PCIe slots. The low profile expansion option provides additional I/O slots 12-16. Slots 12, 13, 15, and 16 are x8 Gen2 PCIe slots. For configuration three, the DL980 with eight VMA arrays requires eight PCIe slots. In order to get the best performance for configurations two and three you will need to also install the low profile expansion option to optimize performance with the PCIe x8 slots. If a VMA Array is ordered with a DL980 G7, PCIe Option I/O Expansion (588137-B21) is integrated regardless. If 3 or more (up to 8) VMA Arrays are ordered with the DL980 G7, the low profile PCIe Option (AM434A) is automatically integrated. Table 6 provides the recommended slot locations occupied by the VMA pass-thru cards for each configuration.

Table 6. HP ProLiant DL980 G7 server with HP VMA PCIe slot configurations

Configuration	Number of HP VMA arrays	DL980 PCIe x8 slots needed	Recommended PCIe x8 I/O slot #	Slot Type
1	2	2	2, 5	x8 Gen2 PCIe
2	4	4	2, 5, 9, 11	x8 Gen2 PCIe
3	8	8	2, 3, 5, 9, 11, 12, 13, 15	x8 Gen2 PCIe

If an add-on SAS controller is installed into the DL980 it may be possible the SAS controller could interfere with the performance of any VMA pass-thru cards installed in PCIe x8 slots 9 and 11 on the Standard Main I/O. You may want to move any VMA pass-thru cards to alternate PCIe x8 slots 6 and 16.

Note:
It is **not supported** to use slot 1 for any HP VMA PCIe Pass Thru Cards due to low I/O performance (PCI-x 100 MHz).

Recommended configurations

The detailed information for these recommended configurations includes the server, number and type of processors, memory, internal and external storage. The configurations were evaluated for a given workload: concurrent users, I/O operations per second, transactions/hour, and database size. The configurations are based on testing done in HP's performance integration lab to support the implementation of online transaction processing servers using HP ProLiant servers running Red Hat Linux and Oracle 11gR2 Enterprise Edition.

The configurations were determined based on the following utilization criteria:

- CPU utilization of approximately 75% at the target workload
- Buffer cache hit ratio of no less than 98% (indicating that the majority of the active database buffers are in physical memory)
- Disk I/O activity reflects a read/write ratio of approximately 80/20

HP ProLiant DL980 G7 server configurations examples:

- 4 or 8 processors
- 256 - 1024 GB of memory (4P upgradeable to 1TB) (8P upgradeable to 2TB)
- 8, 15 and 25 TB database sizes
- 35 million business trans/hr (6,000 users), 70 million business trans/hr (13,000 users) and 150 million trans/hr (26,000 users)
- Storage – 2, 4 or 8 VMA3210 Memory Arrays

Table 7. HP ProLiant DL980 G7 server configurations

Number	CPUs	CPU Type	Memory GB	External Storage	Users	Transactions /hr	Database size TB
1	4	Intel Xeon E7-2830 (24M Cache, 2.13 GHz, 105W, 6.4 GT/s QPI)	256	(2) VMA3210 Memory Arrays	6,000	35,000,000	8
2	4	Xeon E7-4870 (30MB Cache, 2.4GHz, 130W, 6.4 GT/s QPI)	512	(4) VMA3210 Memory Arrays	13,000	70,000,000	15
3	8	Xeon E7-4870 (30MB Cache, 2.4GHz, 130W, 6.4 GT/s QPI)	1024	(8) VMA3210 Memory Arrays	26,000	150,000,000	25

Note:
Currently this reference solution is supported with Red Hat Linux 5.6, with support for Windows Server 2008 R2 shortly after initial release.

Figure 1. HP ProLiant DL980 G7 with eight internal drives and an optical DVD

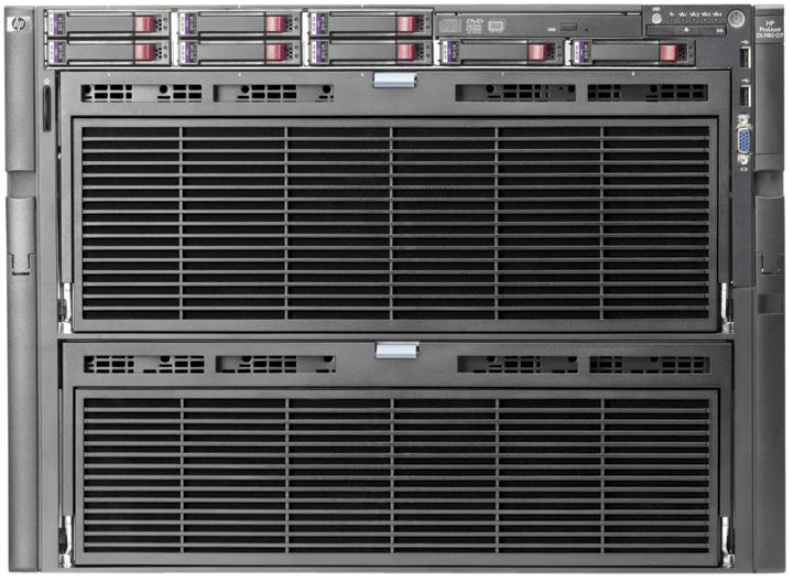

System/Environment setup

Storage configuration details

Internal storage

The DL980 G7 OLTP recommended configurations use four internal drives configured with RAID1 for the OS and 11gR2 software. The server supports up to eight internal drives, so additional drives can be added for staging data, logs, or other requirements.

Table 8. Internal storage controllers

Database Server	Internal Storage Controller
DL980 G7	• HP Smart Array P410i/Zero Memory Controller (RAID 0/1/1+0) • Available upgrades: 512MB with BBWC, Battery kit upgrade (for BBWC), 512MB Flash Backed Write Cache (FBWC), 1G Flash Backed Write Cache, and Smart Array Advanced Pack (SAAP)

External storage

VMA3210 Memory Array

Hewlett Packard's VMA-series Memory Arrays are external systems offering high capacities of solid state storage, using hot-swappable flash memory devices. The VMA-series Memory Arrays feature rackable external solid state storage in capacities ranging from 5TB to 10TB per Single-level Cell (SLC) array. Multiple arrays can create scalable capacities up to petabytes of solid state storage.

The arrays are designed for customers who need low-latency response times, very high I/O performance, and large capacity. Solid state memory technology provides these storage products with no-moving-parts reliability and low power and cooling requirements. These arrays offer a choice of PCIe-direct attach or Fibre Channel SAN attach. PCIe-direct attach with rackmount servers offers

lowest latency speeds for single-server applications. Fibre Channel SAN attach with either blade or rackmount servers allows multiple servers to operate on the same shared data, with practically unlimited amounts of solid state storage.

The VMA3210 Memory Array is a 3U appliance that provides solid state storage to servers. The VMA-series Memory Array eliminates the seek-bound limitations of rotating disk storage and delivers significant application performance benefits over traditional storage systems that utilize hundreds of hard disk drives (HDDs). VMA-series Memory Arrays have practically unlimited scalability – the very large capacities that can be created with VMA-series arrays enable users to place entire databases into solid state memory without needing to architect their solution around limited capacity solid state drives.

Using a massively parallel architecture with distributed garbage collection implemented in hardware, each VMA-series Memory Array with Single Layer Cell (SLC) NAND Flash delivers sustained DRAM-like performance that is an order of magnitude faster than a similar size HDD array:

- 1.4 GB/s bandwidth
- 350K Random Read IOPS
- Delivers up to 240K Sustained Random Write IOPS (4K block with RAID)
- Up to 10TB raw capacity and 6TB Usable flash storage (if configured for best performance)
- Low latency Flash 4 + 1 parity RAID
  - Reads: 80 micro seconds
  - Writes: 25 micro seconds
  - Non-blocking Erases ensures consistent low-latency Reads
- Significant application performance
- 90% reduction in cost, power, and cooling for similar performance
- Power-safe writes without batteries

The configuration used to achieve highest performance consists of attaching the VMA directly to the PCIe slots in the back of the DL980 server using VMA Direct Attach PCIe x8 pass thru cards. This is the configuration used in the three solutions described in this paper. VMA can also be configured via Fibre Channel via a 2U HP VMA SAN Gateway that provides LUN management functions. A maximum of two VMA memory arrays are supported behind one VMA SAN Gateway for Fibre Channel SAN attach. However, latency will be higher than that achieved with direct attach pass thru cards. If an active-active highly available solution is required using Oracle Real Application Clusters (RAC) then you could use a combination of the HP VMA SAN Gateway with the VMAs to provide a shared storage configuration.

Table 9. DL980 G7 and VMA Storage Performance and Capacity

Attribute	VMA Configuration 1	VMA Configuration 2	VMA Configuration 3
Number of arrays	2	4	8
Potential IOPS	600,000	1,200,000	2,400,000
Bandwidth (GB/s)	2.8	5.6	11.2
Capacity (RAW)	20TB	40TB	80TB
Capacity (Useable)	12TB	24TB	48TB

When the active data of large-scale and high performance applications using Oracle databases is assigned to VMA appliances, dramatic application acceleration and much higher CPU utilization is achieved. This enables the data center by supporting high performance random I/O without the need for large DRAM server footprints or thousands of spinning HDDs.

Each VMA3210 is populated with 84 x 128GB solid state memory (SSM) Modules based on Single Layer Cell (SLC) flash memory devices. VMA Switched Memory technology and VMA's Flash RAID provide the system with its industry-leading scalability, data reliability and efficiency:

- Hot swappable flash memory modules with 4+1 RAID and 4 spares
- System power: less than 120W per Useable Terabyte
- Endurance: greater than 8TB Writes per Hour for 5 Years
- Redundant power supplies and cooling are provided in all VMA-series Memory Array systems.

Manageability is built into the system with an on board management controller.

- Built-in management processor with command line interface
- Front panel status icons and internal status LEDs on each memory board
- Appliance with factory-installed VMA controller software.

The optional FC SAN Gateway appliance and vShare software supports a web-based user interface.

Figure 2 shows a single VMA-series Memory Array.

Figure 2. VMA3210 Memory Array

Figure 3 depicts the direct connectivity to the server.

Figure 3. Depicts the direct connections between the VMA units and the DL980 server

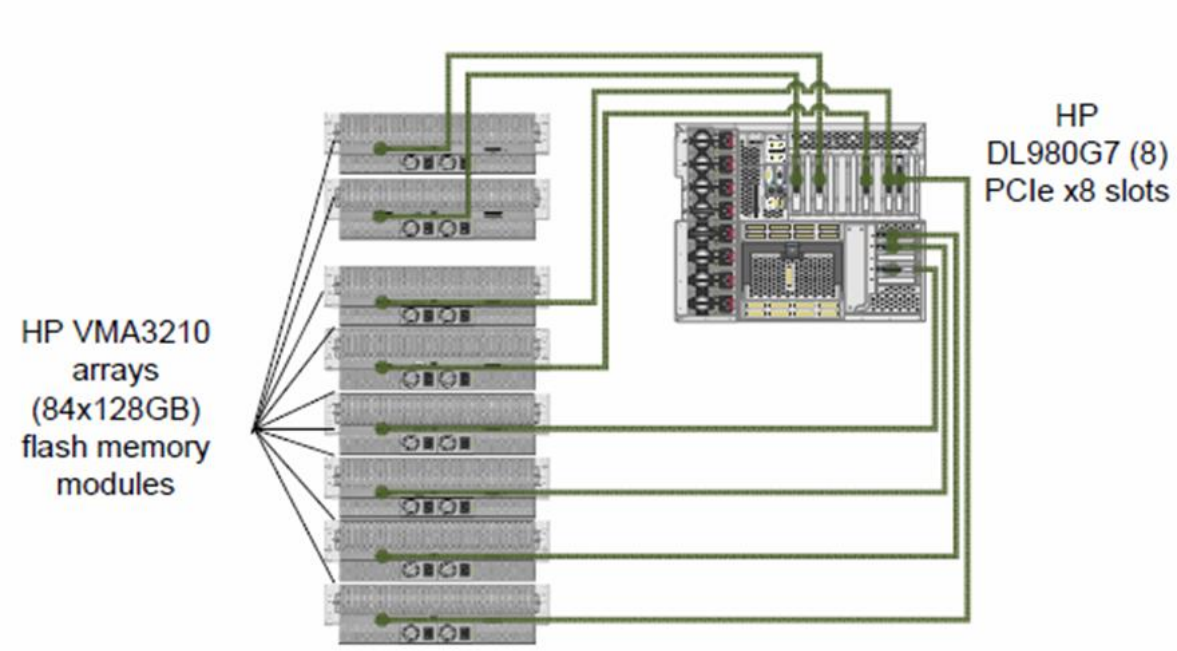

Note:
Factory Integration and racking is supported for DL980 servers ordered with VMA-series Memory Array and PCIe-direct options.

HP Care Pack Services: Packaged server and storage services for increased uptime, productivity and ROI.

Additional external storage requirements

A typical data center will have some sort of shared conventional storage for the various applications based on performance needs and business storage requirements. Most business solutions will have requirements for multiple environments besides the primary production storage requirements such as development, test, quality, training, sandbox and backup. Each of these additional environments typically will require their own storage solution. Many of these non-production environments do not require the extreme performance of the production solutions and therefore can be hosted on lower cost, higher capacity storage solutions. The HP P6000/Enterprise Virtual Array (EVA) disk array could make a great backup solution for the VMA units.

There are at least three options to hosting the storage for these other environment requirements:

- Use the same storage architecture as production
- Use existing conventional storage in the current data center such as a storage area network (SAN)
- Add additional disk array storage

This reference architecture includes additional conventional storage recommendations for these multiple non-production environments. However, you can substitute an existing SAN solution or add additional VMA arrays to satisfy these non-production requirements.

HP P6300 EVA – The HP P6000 Enterprise Virtual Array (EVA) is an enterprise class virtual storage array family for midsized customers at an affordable price. With built in virtualization, the P6000 EVA is designed to improve capacity utilization and be easy to manage, which lowers the cost of ownership compared to traditional arrays. These arrays have high performance, scale easily, and are highly reliable and available. The P6000 EVA is a trusted platform for enterprise application consolidation with solutions for Microsoft Exchange and SQL Server, Oracle, and SAP. It also is a great platform for all non-production environments.

The HP P6000 EVA family is the next generation of EVA with two new models: P6300 EVA and P6500 EVA. The P6000 EVA family supports Serial Attached SCSI (SAS) disk drives with a wide range of Small Form Factor (SFF) SAS drives, Large Form Factor (LFF) SAS drives, or combinations of both, to better match customer storage capacity, performance, power, and availability needs. The P6000 EVA models are available with Fibre Channel host ports only, or Fibre Channel and 1Gb/s iSCSI host ports, or Fibre Channel and 10Gb/s iSCSI/FCoE host ports. P6000 EVA is also very energy efficient, with high efficiency power supplies for both the controllers and disk drive enclosures, with temperature sensing, self-adjusting variable speed fans.

HP P6000 Command View management software now includes Thin Provisioning, which helps reduce the storage capacity required, to help reduce the number of disks needed, and thus reduce power and cooling costs. P6000 EVA also offers robust local and remote replication capabilities with HP P6000 Business Copy and HP P6000 Continuous Access software. P6000 Business Copy now also supports Dynamic LUN and RAID migration, where in one step the user can change the characteristics of an existing LUN, such as size, RAID type, or disk type while the host I/O operation is active. This feature is useful, for example, to move data to more space efficient storage tiers or to move a heavily accessed LUN to a different disk group to improve performance.

The reference configurations deploy two disk groups within the EVA. For example, the test/development database volumes are in one disk group and the archive logs, flashback recovery, and RMAN backup volumes are in the second disk group. HP 6300 Enterprise Virtual Arrays are configured and managed using the HP Command View EVA Software.

Note:
Always check HP's Single Point of Connectivity Knowledge (SPOCK) compatibility and interoperability matrix for a "configuration set" of Storage Area Network (SAN) components. It is also very important to examine the notes in each section, because they may further limit or clarify supported configurations.

See http://www.hp.com/storage/spock (requires an HP Passport account).

Figure 4 shows a sample drive configuration for the P6300 EVA.

Figure 4. P6300 EVA logical drive configuration for the additional nonproduction environments

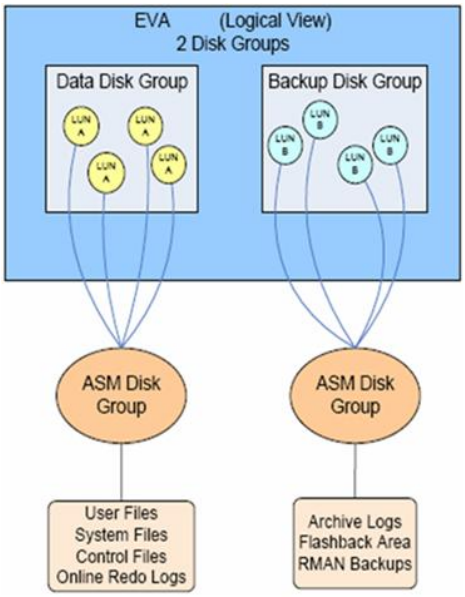

DL980 system setup

When setting up the DL980 for an Oracle database there are a few best practices that will provide improved performance and proper configuration.

- If the DL980 G7 has 1.0TB or more of RAM, the BIOS System Address Mode should be set to 44-bit. By default this is set to 40-bit.
- Spread the same size DIMM memory evenly across all memory cartridge sockets for maximum performance. For the two 4-socket configurations it is best to populate all memory slots associated with those four CPUs. For the 8-socket configuration install memory DIMMs across all memory sockets of the eight CPUs for optimum NUMA performance.
- Use only dual or quad rank DIMMs as they are faster than single rank DIMMs.
- Configure per DL980 best practices to minimize memory latencies.
- If the application has parallelism that will work optimally with Hyper-Threading, leaving it enabled in the BIOS could result in significance improvement. If not, Hyper-Threading can cause a drop in overall performance. We saw up to 30% improvement in our testing by turning Hyper-Threading on. Make sure to test it for your particular workload.
- Use the PCI-Express x8 slots for critical disk I/O.
- If connecting three or more HP VMAs, for optimal performance install the Low Profile PCIe expansion option.
- Distribute PCIe pass-thru cards evenly across the available I/O bays and I/O hubs.
- Install and configure the NUMA packages.

- Enable the Hardware Prefetcher. The Hardware Prefetcher proactively collects data from the memory to the cache in order to reduce response time. If the application is memory intensive, this enablement may have a negative impact on the performance.
- In the BIOS menu under Power Management Options set the HP Power Profile to Maximum Performance.
- For environments not running Virtual Machines – Disable Virtualization support in the BIOS as well as INTEL VT-D.
- By default Drive Write Cache is Disable and should remain disabled using VMA attached storage.
- Do not use slot 1 for any HP VMA PCIe Pass Thru Cards. It is a PCI-x Gen1 slot.
- Set HUGE pages for system global area (SGA) to 2MB. The main advantages of creating an SGA using huge pages has to do with increased performance from improving the translation lookaside buffer (TLB) hit ratio and reducing the memory footprint required for mapping the SGA.

Oracle database setup

Location of Oracle 11gR2 binaries

The Oracle database software should be installed separately from the storage used for the database instance itself. The recommended configurations are designed with the intent that the Oracle binaries be placed on the same drives as the OS. There is no performance benefit to separating the OS and Oracle binaries on different drives, though some customers will chose to do so to simplify system backup and maintenance operations.

Oracle databases location

For the HP VMA storage configuration it is recommended that the Oracle database components, such as data, indexes, undo, temp and redo, should be managed with Oracle Automatic Storage Management (ASM) to stripe across the storage arrays. ASM should be used in "external redundancy" mode, since the VMA arrays themselves will provide the RAID level protection for the databases.

The HP P6300 EVA is used for additional database storage requirements such as database backups, flash recovery area, storing archive logs, test, QA, training, sandbox and development environments. There is fibre channel connectivity between the P6300 EVA and the production and test/dev DL980 servers through a pair of fibre channel switches.

Memory allocation for OS

In an Oracle 11gR2 setup it is recommended to use any extra memory available on the system for the system global area (SGA). This can improve I/O performance. Leave 10% of the memory available for the operating system and about 4% for buffer cache.

Patching

Make sure the Linux operating system and Oracle database is patched to current patch releases for best performance and problem resolution. Please refer to the Oracle database installation guide and release notes for recommended patch, OS, and database parameter recommendations.

Other best practice recommendations you may want to consider:

- Enable NUMA mode by setting the Oracle database parameter:
  - `_enable_NUMA_support = true`
  - If NUMA is enabled, alert.log will trace it as *"NUMA system found and support enabled"*

- We recommend disabling automatic memory management:
  - `SGA_TARGET = 0`
  - `MEMORY_TARGET = 0` (Oracle DB 11g)
  - `MEMORY_MAX_TARGET = 0` (Oracle DB 11g)
- Set automatic memory management for the program global area (PGA):
  - `PGA_AGGREGATE_TARGET` and `WORK_AREA_POLICY` set to `AUTO`

VMA-series Memory Array storage setup

VMA arrays can be connected directly to existing hardware or can be combined into a hardware refresh or new opportunities. The following best practices are recommended when using the DL980 and VMA in an Oracle database environment:

- The first step towards installing the VMA drivers is to ensure that all of the required VMA packages are installed on the host server prior to the installation and configuration of the VMA driver.
- The VMA arrays are pre-formatted at 65% of the available storage for improved write performance.
- Slots capable of bi-directional x8 PCIe provide the best performance.
- On the server enable asynchronous I/O.
- Flash-based memory arrays are designed for 4KB block access or any multiple of 4KB blocks. Smaller block sizes (for example, 512 bytes) will significantly reduce performance, particularly for writes. Therefore, it is import that you verify whether the file system and operating system are 4KB aligned.
- Use the Parted software tool to create 4 to 16 partitions across each array and leverage ASM to concatenate the volumes into a common disk group or groups.
  - Be sure to set the offset in Parted to 1MB to ensure a 4KB boundary for I/O.
  - 512B partitions start at 2048s
  - 4096B partitions start at 256s
- Balance the workload evenly across all VMA Arrays.
  - Depending on your workload set the allocation unit (AU) to 1MB or 4MB. When using Oracle Automatic Storage Management (ASM), you have an option to enable software mirroring. Our recommendation is not to implement Oracle ASM mirroring because the arrays are designed to meet five 9s of high availability without incurring the I/O overhead of Oracle mirroring or the 50% reduction in usable storage space.
- Recommended to increase the I/O Q-depth to 256 on the DL980 (default is 16).
- It is best to use dual 208V power cords/supplies for most optimal performance.
- For applications that require high availability, the use of Oracle Real Application Clusters (RAC) is fully supported; however, it is not required. For customers that are using Oracle Enterprise Edition, Oracle Data Guard configured with Fast Failover is the recommended high availability solution. This solution is a good match for the simplicity of the HP VMA configuration. There are also other replication products that are fully supported with VMA arrays.
- HP and Oracle both recommend using the Linux kernel's Huge Pages feature (which necessitates the disabling of Oracle Automatic Memory Management).

Test assumptions

- These configurations are examples; there are many additional server/storage options which could also meet the workload requirements.
- Testing has not been performed to date on each of these exact configurations. The configurations are based on best practices, benchmarks, extrapolation based on current test knowledge, and performance assumptions through discussions with Oracle and HP experts.
- No consideration at this point has been given for spares, recovery storage groups, clusters, etc.; nevertheless, all of these options can be added to the base configurations.

Note:
One PCIe Pass Thru Card is required with each Memory Array and is used to connect the array to the external server. The card must be installed in the target server (for the PCI-direct attach method)

Bill of materials

Recommended configuration 1 for 6,000 users

Figure 5. Oracle Database 11gR2 online transaction processing (OLTP) 6000 user solution

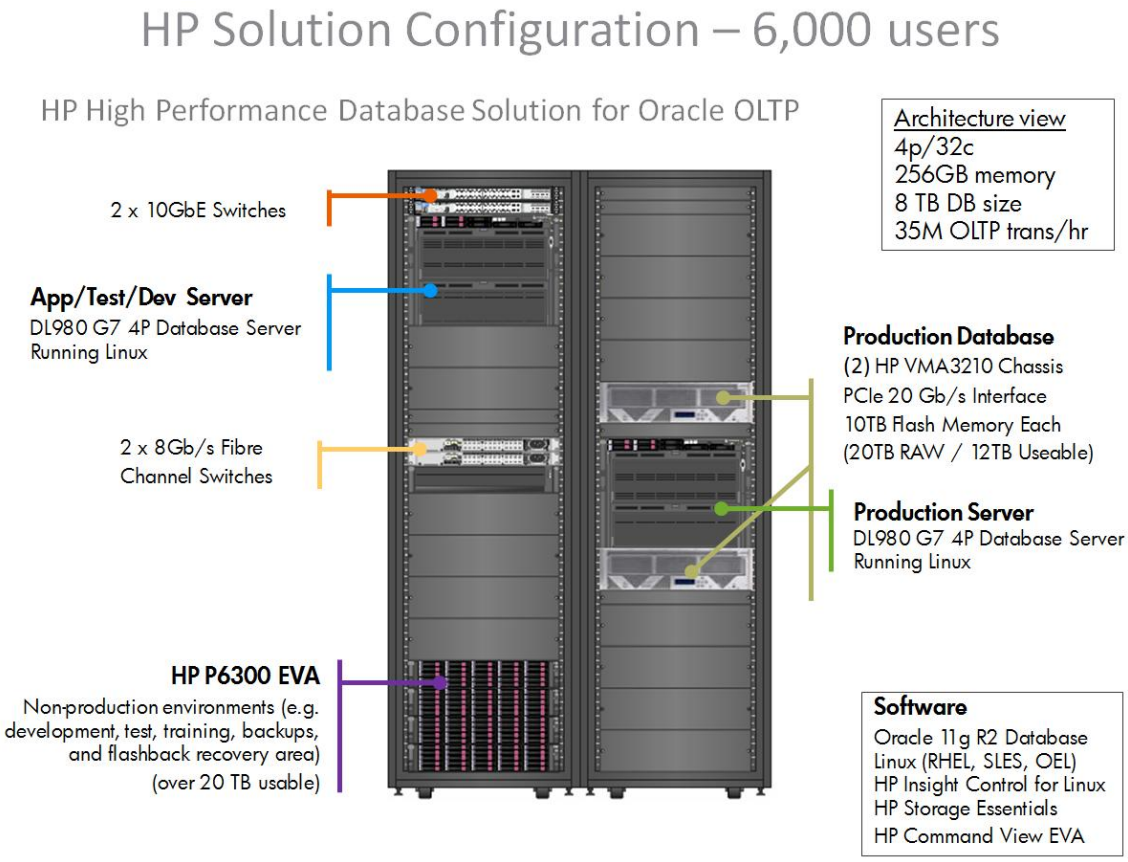

Table 10. Bill of materials – Configuration 1 – DL980 G7 4P with (2) VMA3210 Memory Arrays

Qty	Description
	Production Database Server Configuration
1	HP ProLiant DL980 G7 CTO system
1	HP DL980 G7 E7-2830 FIO 4-processor Kit
4	HP 1200W CS Platinum Power Supply kit
32	HP 8GB 2Rx4 PC3-10600R-9 kit
1	HP Slim 12.7mm SATA DVD Optical kit
1	PCI Express I/O Exp. Kit
2	HP 8Gb Dual Port PCIe FC HBA
1	Dual Port 10GbE Server Adapter
4	HP 72GB 6G SAS 15k 2.5in DP ENT HDD
	Test/Dev Database Server Configuration
1	HP ProLiant DL980 G7 CTO system
1	HP DL980 G7 E7-2830 FIO 4-processor Kit
4	HP 1200W CS Platinum Power Supply kit
16	HP 8GB 2Rx4 PC3-10600R-9 kit
1	HP Slim 12.7mm SATA DVD Optical kit
1	PCI Express I/O Exp. Kit
2	Dual port 8 Gb FC HBA
1	Dual Port 10GbE Server Adapter
4	HP 72GB 6G SAS 15k 2.5in DP ENT HDD
	Storage Configuration
2	HP VMA3210 10TB SLC Memory Array with PCIe 20 Gbit/s Interface
2	HP VMA PCIe Pass Thru Card, Dual Port PCIe x8 Interfaces
2	24 port 8Gb/s Fibre Channel Switches
1	HP P6300 Enterprise Virtual Array dual controller with dual embedded switch option with 86 300GB 10K rpm SFF 6Gb/s SAS HDD
	Other
2	24 port 10GbE Switches

Recommended configuration 2 for 13,000 users

Figure 6. Oracle Database 11gR2 online transaction processing (OLTP) 13,000 user solution

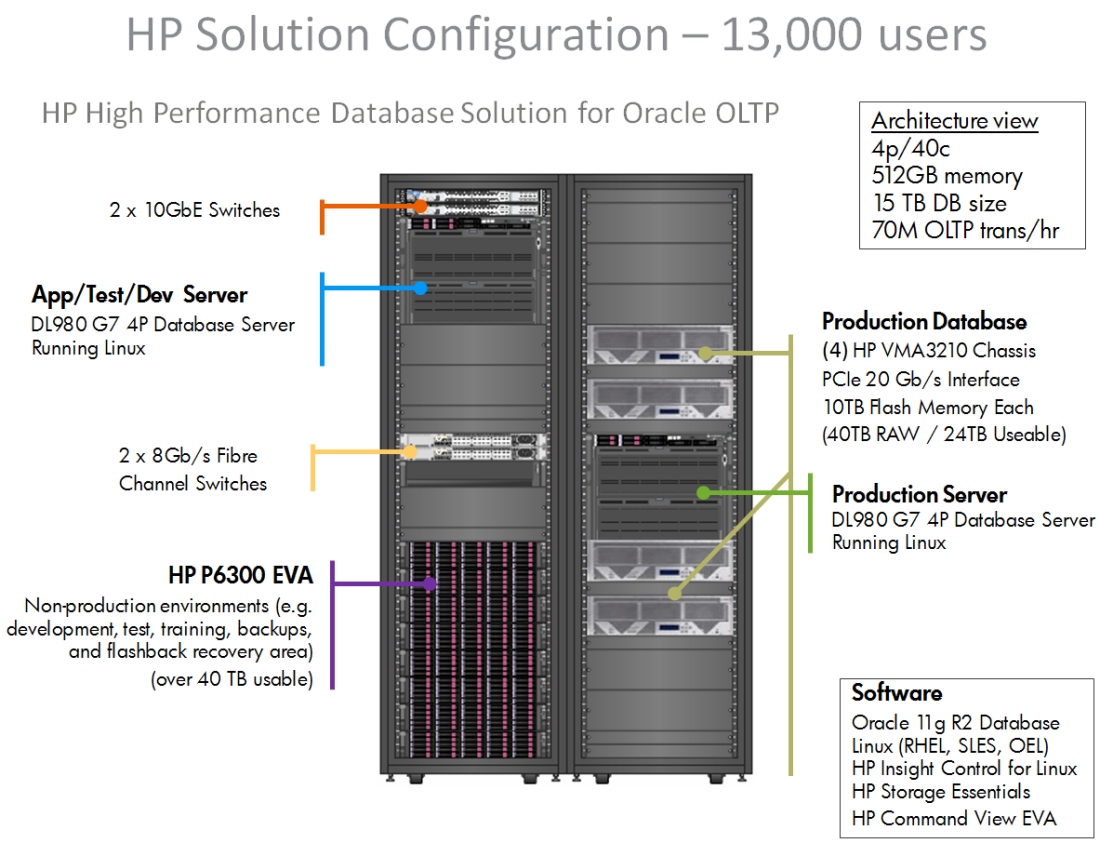

Table 11. Bill of materials – Configuration 2 – DL980 G7 4P with (4) VMA3210 Memory Arrays

Qty	Description
	Production Database Server Configuration
1	HP ProLiant DL980 G7 CTO system
1	HP DL980 G7 E7-4870 FIO 4-processor Kit
4	HP 1200W CS Platinum Power Supply kit
64	HP 8GB 2Rx4 PC3-10600R-9 kit
1	HP Slim 12.7mm SATA DVD Optical kit
1	PCI Express I/O Exp. Kit
2	HP 8Gb Dual Port PCIe FC HBA
1	Dual Port 10GbE Server Adapter

Qty	Description
4	HP 72GB 6G SAS 15k 2.5in DP ENT HDD
	Test/Dev Database Server Configuration
1	HP ProLiant DL980 G7 CTO system
1	HP DL980 G7 E7-2830 FIO 4-processor Kit
4	HP 1200W CS Platinum Power Supply kit
16	HP 8GB 2Rx4 PC3-10600R-9 kit
1	HP Slim 12.7mm SATA DVD Optical kit
1	PCI Express I/O Exp. Kit
2	Dual port 8 Gb FC HBA
1	Dual Port 10GbE Server Adapter
4	HP 72GB 6G SAS 15k 2.5in DP ENT HDD
	Storage Configuration
4	HP VMA3210 10TB SLC Memory Array with PCIe 20 Gbit/s Interface
4	HP VMA PCIe Pass Thru Card, Dual Port PCIe x8 Interfaces
2	24 port 8Gb/s Fibre Channel Switches
1	HP P6300 Enterprise Virtual Array dual controller with dual embedded switch option with 170 300GB 10K rpm SFF 6Gb/s SAS HDD
	Other
2	24 port 10GbE Switches

Figure 7. Oracle Database 11gR2 online transaction processing (OLTP) 26,000 user solution

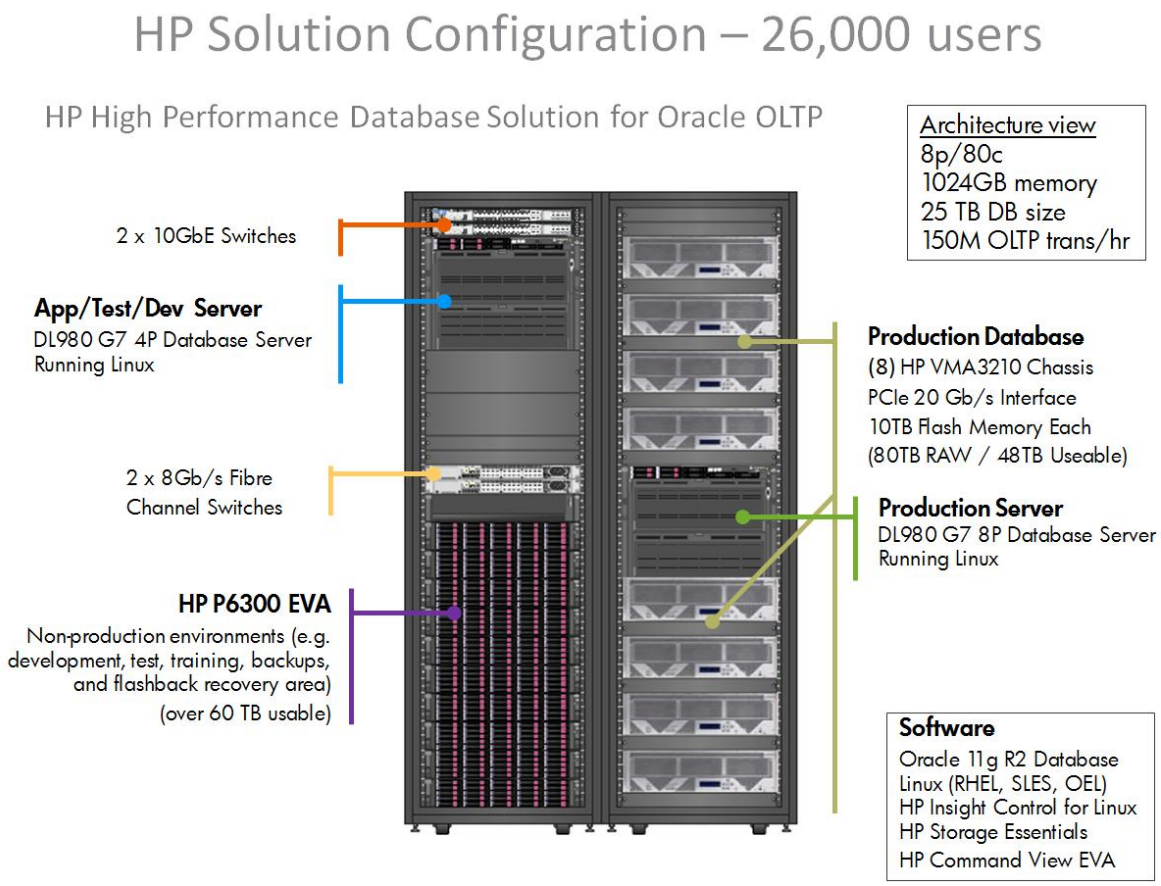

Table 12. Bill of materials – Configuration 3 – DL980 G7 8P with (8) VMA3210 Memory Arrays

Qty	Description
	Production Database Server Configuration
1	HP ProLiant DL980 G7 CTO system
2	HP DL980 G7 E7-4870 FIO 4-processor Kit
1	HP DL980 CPU Installation Assembly
4	HP 1200W CS Platinum Power Supply kit
128	HP 8GB 2Rx4 PC3-10600R-9 kit
8	HP DL980 G7 Memory Board
1	HP Slim 12.7mm SATA DVD Optical kit

Qty	Description
1	PCI Express I/O Exp. Kit
1	Low Profile PCI Express I/O Expansion Kit
2	HP 8Gb Dual Port PCIe FC HBA
1	Dual Port 10GbE Server Adapter
4	HP 72GB 6G SAS 15k 2.5in DP ENT HDD
	Test/Dev Database Server Configuration
1	HP ProLiant DL980 G7 CTO system
1	HP DL980 G7 E7-2830 FIO 4-processor Kit
4	HP 1200W CS Platinum Power Supply kit
16	HP 8GB 2Rx4 PC3-10600R-9 kit
1	HP Slim 12.7mm SATA DVD Optical kit
1	PCI Express I/O Exp. Kit
2	Dual port 8 Gb FC HBA
1	Dual Port 10GbE Server Adapter
4	HP 72GB 6G SAS 15k 2.5in DP ENT HDD
	Storage Configuration
8	HP VMA3210 10TB SLC Memory Array with PCIe 20 Gbit/s Interface
8	HP VMA PCIe Pass Thru Card, Dual Port PCIe x8 Interfaces
2	24 port 8Gb/s Fibre Channel Switches
1	HP P6300 Enterprise Virtual Array dual controller with dual embedded switch option with 250 300GB 10K rpm SFF 6Gb/s SAS HDD
	Other
2	24 port 10GbE Switches

Implementing a proof-of-concept

As a matter of best practice for all deployments, HP recommends implementing a proof-of-concept using a test environment that matches as closely as possible the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HP Sales representative or your HP partner.

Appendix

Oracle Advanced Compression

Many customers are looking for solutions that provide a means for reducing the size of their rapidly growing databases without negatively affecting their end user performance. Oracle 11gR2 offers integrated database compression to address this requirement.

We often think of compression as being a trade-off between performance and storage: compression reduces the amount of storage required, but the overhead of compressing and decompressing makes things slower. However, while there is always some CPU overhead involved in compression the effect on table scan I/O can be favorable, since if a table is reduced in size it will require fewer I/O operations to read it.

Prior to 11g, table compression could only be achieved when the table was created, rebuilt or when using direct load operations. However, in 11gR2, the Advanced Compression option allows data to be compressed when manipulated by standard DML (Data Manipulation Language). The data compression feature in Oracle 11gR2 Enterprise Edition reduces the size of tables and indexes while providing full row level locking for updates. There are two types of compression.

1. Row compression enables storing fixed-length data types in a variable-length storage format.
2. Page compression is a superset of row compression. It minimizes the storage of redundant data on the page by storing commonly-occurring byte patterns on the page once, and then referencing these values for respective columns.

Oracle's Advanced Compression offers three distinct levels: low, medium, and high. HP and Oracle recommend using the "low" method for best overall OLTP workload performance when data compression is desired. Oracle has provided a compression algorithm specifically designed to work with OLTP type workloads. This recommendation is based upon tests performed by HP and Oracle on industrial-standard x86 hardware (see the reference at the end of this document). Users may wish to evaluate other compression options to determine if the "medium" or "high" setting offers superior performance for their specific workload.

As one would expect, Oracle Advanced Data Compression was very effective at reducing disk utilization of traditional storage arrays. The result was improved data transfer from storage into the database instance for processing and reduced I/O wait overhead. Testing conducted by HP's Oracle Alliances team showed that Advanced Data Compression scaled linearly across the full range of CPU cores on HP 8-socket servers. All indications are that data compression will have an even greater positive impact on performance with the VMA series Memory Array solutions.

For more information

For additional Oracle solutions from HP, please visit http://www.hp.com/go/oracle

HP VMA-series Memory Array products at www.hp.com/go/VMA; www.hp.com/go/vma-docs

HP ProLiant servers, http://www.hp.com/go/proliant

HP P6000 Enterprise Virtual Array (EVA), http://www.hp.com/go/eva

HP P6300/P6500 Enterprise Virtual Array Best Practice
http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA3-2641ENW

HP Single Point of Connectivity Knowledge (SPOCK) website
http://h20272.www2.hp.com

Open Source and Linux from HP www.hp.com/go/linux

11gR2 Compression Tests using Oracle Enterprise Linux,
http://hporacle.com/-/documents/view.asp?id=765
requires registration at http://hporacle.com

Oracle Database Compression with HP DL785 and EVA: a scalability study,
http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA1-0234ENW&cc=us&lc=en

A complete list of certified third party storage for HP servers can be found at:
www.hp.com/products1/serverconnectivity/mass_storage_devices.html

For assistance with HP VMA Array support and training please send an email to:
hpsupport@vmem.com

Learn more: www.hp.com/services/servers and www.hp.com/services/storage

If you do not know your HP sales or HP partner representative, please do not hesitate to send your response via e-mail to one of the regional *HP Oracle Solution Teams* below.

- Americas ssa@hp.com
- Europe / Middle East / Asia oracle-sizing.emea@hp.com
- Asia Pacific / Japan oracle-sizing.apj@hp.com
- Mexico oracle-sizing.mexico@hp.com
- Brasil oracle-sizing.brasil@hp.com
- Other Latin America / Caribbean oracle-sizing.lac@hp.com

To help us improve our documents, please provide feedback at
http://h71019.www7.hp.com/ActiveAnswers/us/en/solutions/technical_tools_feedback.html.

4AA3-2367ENW, Created December 2010; Updated September 2011, Rev. 4